arXiv:1407.4578v1 [stat.ME] 17 Jul 2014

Maximal Autocorrelation Functions in Functional Data Analysis

Giles Hooker* and Steven Roberts†

Abstract

This paper proposes a new factor rotation for the context of functional principal components analysis. This rotation seeks to re-represent a functional subspace in terms of directions of decreasing smoothness as represented by a generalized smoothing metric. The rotation can be implemented simply and we show on two examples that this rotation can improve the interpretability of the leading components.

Keywords: Factor Rotation; Functional Data; Interpretability; Principal Components Analysis

1 Introduction

This paper proposes a new factor rotation for functional principal components analysis (fPCA). In functional data analysis, the use of principal components has received considerable attention; means of defining principal components were studied in Rice and Silverman (1991) and Silverman (1996) and for sparsely observed curves in Yao et al. (2005a); Peng and Paul (2009) and Paul and Peng (2011). Following carrying out fPCA, one of the most common means of dealing with functional covariates has been by employing principal components scores within multivariate methods. Particular examples include linear models (Yao et al., 2005b; Hall et al., 2006; Goldsmith et al., 2011), as responses (Chiou et al., 2004; Sentürk and Müller, 2006) and additive models (Müller and Yao, 2008). In some of these cases, interpretation is gained by combining coefficients from the multivariate model with the principal component functions to create a functional parameter – see the functional linear models in Yao et al. (2005b) – but this is not always possible, as in the additive models in Müller and Yao (2008), in which case the model must be interpreted by treating the principal component directions as having particular meanings.

Despite this interest in fPCA, little has been proposed by the way of factor rotations that might make principal components directions more interpretable. Ramsay and Silverman (2005) examine an extension of the VARIMAX rotation (Kaiser,

*Department of Biological Statistics and Computational Biology, Cornell University, gjh27@cornell.edu

†Research School of Finance, Actuarial Studies and Applied Statistics, Australian National University, steven.roberts@anu.edu.au

1958) from multivariate factor analysis which has the tendency to produce components that focus on particular ranges of the domain of the functions. Liu et al. (2012) propose a rotation towards periodic components in a remote sensing example with functions that cover multiple years with a distinct annual signal. Other methods from the multivariate factor rotation literature could be considered, but we have found no other suggested factor rotations that make use of the structure of functional data. In this paper, we propose a rotation towards maximally smooth principal components. These are the directions in which there is greatest predictability over time and which are also more interpretable.

The factor rotation that we propose is derived from the definition of Min/Max Autocorrelation factors (MAF) introduced by Shapiro and Switzer (1989) and Switzer and Green (1984) for the analysis of gridded multivariate data and parallel time series data, respectively. The principal underlying MAF is to find linear combinations of the original data that have maximum autocorrelation. This property of MAF is in contrast to PCA which finds linear combinations that have maximum variance. Of particular importance for our setting is the fact that when applied to parallel time series a MAF analysis finds linear combinations of the data that are decreasingly smooth functions of time or, in other words, that are decreasingly predictable functions of time. In this vein, we show that a functional analogue of MAF can be obtained that searches for the rotated components that have smallest integrated first derivative. We then demonstrate how this can be extended to any notion of smoothness as given by a linear differential operator defined in Ramsay and Silverman (2005). In our examples, we have developed our methods based on the numerical machinery in the `fda` package in `R` (see Ramsay et al., 2009) but they can be readily employed with alternative functional data representations.

In recent literature, interest in the application and theoretical properties of MAF has been increasing (Cunningham and Ghahramani, 2014; Gallagher et al., 2014). Of particular relevance to the current work is the paper of Henderson et al. (2009) that compares PCA and MAF in the context of ToF-SIMS image interpretation. The authors conclude that MAF is more effective than PCA for the analysis of high signal intensity data. The importance of MAF for forecasting has been investigated by Woillez et al. (2009) in the context of fish stocks. The utility of MAF for forecasting correlated functional data is ongoing work.

The remainder of the paper is structured as follows: we derive the notion of maximally smooth rotation as an analogue of maximally autocorrelated time series in Section 2. The extension of these to any linear differential operator and its implementation in basis expansion methods is given in Section 3 and we demonstrate the effect of these methods in Section 4. We finish with some concluding remarks and further directions.

2 Maximal Autocorrelation Factor Rotations (MAFR)

Our methods are developed on top of the Maximal Autcorrelation Functions proposed in Switzer and Green (1984) for multivariate time series. Suppose that we have a multivariate time series $\mathbf{x}_1, \ldots, \mathbf{x}_T$ in which at each time point the observed

$\mathbf{x}_t$ is a vector. The maximally autocorrelated time series is the linear transformation $\mathbf{b}$ such that $\text{cor}(\mathbf{x}_t\mathbf{b}, \mathbf{x}_{t+1}\mathbf{b})$ is maximized. In order to apply this to functional data analysis, we re-interpret the criterion as

$$\hat{\mathbf{b}} = \underset{\mathbf{b}}{\text{argmin}} \frac{\sum_{t=1}^{T-1} \left(\mathbf{b}^T\mathbf{x}_{t+1} - \mathbf{b}\mathbf{x}_t\right)^2}{\mathbf{b}^T \sum_{t=1}^{T} \mathbf{x}_t\mathbf{x}_t^T\mathbf{b}}.$$

In a functional data analysis context we consider $\mathbf{x}_t$ to derive from the evaluation of a vector of functions $\mathbf{x}(t)$ at times $t = i(\Delta t)$ for $i = 1, \ldots, T$. By dividing by Δt we can re-represent this criterion as

$$\hat{\mathbf{b}} = \underset{\mathbf{b}}{\text{argmin}} \frac{\int \mathbf{b}^T\dot{\mathbf{x}}(t)\dot{\mathbf{x}}(t)^T\mathbf{b}dt}{\mathbf{b}^T \int \mathbf{x}(t)\mathbf{x}(t)^T dt\mathbf{b}}.$$

where $\dot{\mathbf{x}}(t)$ is the vector of time-derivatives of $\mathbf{x}(t)$.

Here, we recognize the numerator as having the form of a classical first-derivative smoothing penalty on the univariate function $z(t) = \mathbf{b}^T\mathbf{x}(t)$. In this spirit, we can more generally define a criterion by any linear differential smoothing operator L as in Ramsay and Silverman (2005). This allows us to define the MAFR criterion as

$$\text{MAFR}_L(\mathbf{b}) = \frac{\int \mathbf{b}^T L\mathbf{x}(t) L\mathbf{x}(t)^T\mathbf{b}dt}{\mathbf{b}^T \int \mathbf{x}(t)\mathbf{x}(t)^T dt\mathbf{b}}$$

where the operator L is a linear combination of derivatives:

$$L\mathbf{x}(t) = \frac{d^k}{dt^k}\mathbf{x}_t + \sum_{j=0}^{k=1} a_j(t)\frac{d^j}{dt^j}\mathbf{x}(t).$$

The most common choices for L correspond to the first and second derivatives: $L\mathbf{x}(t) = \dot{\mathbf{x}}(t)$ or $L\mathbf{x}(t) = \ddot{\mathbf{x}}(t)$ but more complex penalties can also be useful and we examine the *harmonic acceleration* penalty

$$L\mathbf{x}(t) = \frac{d^3}{dt^3}\mathbf{x}(t) - \frac{\omega}{2\pi}\frac{d}{dt}\mathbf{x}(t)$$

which defines smoothness in terms of sin and cosin functions with period ω as well as constant shifts (see Ramsay and Silverman, 2005).

We have written our criterion in terms of a collection of functions $\mathbf{x}(t)$ above, but this method is treated as a factor rotation to be applied following fPCA with a fixed number of components selected. Thus, below we will replace $\mathbf{x}$ with $\boldsymbol{\phi}(t) = (\phi_1(t), \ldots, \phi_K(t))$ to conform to more common notation. If the dimension of $\mathbf{x}(t)$ is allowed to grow, we will always be able to reduce MAFR by adding further components, yielding rotations in which $L\mathbf{b}\mathbf{x} \to 0$. This same phenomenon occurs for classical factor rotations in multivariate analysis when the number of variables increases or in MAFs with an increasing number of time series. Liu et al. (2012) found that applying VARIMAX rotations to a large number of principal components resulted in essentially uninterpretable results. Similar comments may be made

about the maximal autocorrelation functions in Switzer and Green (1984). Along similar lines, we expect that the trailing components after rotation will be the least interpretable and there is a trade-off between increasing the smoothness of the leading components and allowing some variation to be absorbed into the remaining, less-interpretable versions. These comments also apply to other factor rotations, although in the examples below we find that the leading components are smoothed while the remaining ones are relatively unaffected.

3 Numerical Implementation

In this section, we describe the numerical implementation of the factor rotation. This can be accomplished easily using the basis expansion methods in the `fda` package in `R` (Ramsay et al., 2013, 2009), but it relies only on our ability to obtain inner products of the derivatives of principal component functions.

We assume that a set of principal components $\boldsymbol{\phi}(t)$ have been obtained from data. Since these are orthonormal by definition, we have $\int \boldsymbol{\phi}(t)^T \boldsymbol{\phi}(t) dt$ is given by the identiy, and thus the MAFR rotation corresponds to

$$\hat{\mathbf{b}} = \underset{\mathbf{b}}{\operatorname{argmin}}\, \mathbf{b}^T \left[\int L\boldsymbol{\phi}(t)^T L\boldsymbol{\phi}(t) dt \right] \mathbf{b}, \text{ subject to } \mathbf{b}^T\mathbf{b} = 1$$

By standard arguments, the solution to this problem is the smallest eigenvector of the matrix

$$P = \left[\int L\boldsymbol{\phi}(t)^T L\boldsymbol{\phi}(t) dt \right].$$

We may define successive rotations $\mathbf{b}_2, \ldots, \mathbf{b}_k$ by minimizing MAFR($\mathbf{b}$) subject to $\mathbf{b}_i^T \mathbf{b}_j = I_{i=j}$. These are given by the succeeding columns of U in the Eigen-decomposition

$$P = UDU^T.$$

We can thus define a rotation to new components

$$\boldsymbol{\psi}(t) = U^T \boldsymbol{\phi}(t)$$

If, as is standard, the diagonal matrix D is ordered from largest to smallest eigenvalues, the final components of $\boldsymbol{\psi}$ should be the smoothest. We observe that since both U and the $\boldsymbol{\phi}$ are orthonormal, so are the $\boldsymbol{\psi}$.

If we have an original set of curves represented in terms of principal component scores

$$x_i(t) = \sum_{k=1}^{K} s_{ij}\phi_j(t) = \mathbf{s}_i^T \boldsymbol{\phi}(t)$$

the score vector $\mathbf{s}_i$ can be re-represented in the basis defined by $\boldsymbol{\psi}(t)$ in terms of $\mathbf{t}_i = U^T \mathbf{s}_i$. If the variances of the original retained principal components are given in the diagonal matrix Σ, the MAFR components have scores with associated covariance $U^T \Sigma U$.

4 Examples

4.1 A Simulated Experiment

We begin by experimenting with the effect of this rotation on simulated data in which rotation should help to capture a "true" set of leading principal component directions. For this simulation we represented 100 curves via a Fourier basis expansion with bases

$$f_0(t) = 1,\ f_{2i}(t) = \sin(2\pi i t),\ f_{2i+1}(t) = \cos(2\pi i t)$$

and simulated the coefficients of the first 25 such basis functions as independent normals with exponentially decreasing variance:

$$x_i(t) = \sum_{j=0}^{24} c_j f_j(t),\ c_j \sim N(0, \exp(-j/4)).$$

Under this framework, 10 principal components are required to capture 99% of the variation in these data.

We employed a rotation based on minimizing the harmonic acceleration of the leading components and have plotted the original curves along with the first four and final two components in Figure 1. Here we see there is a distinct smoothing of the leading components. Interestingly, the final MAFR component is more purely sinusoidal – with a higher frequency than the harmonic acceleration penalty – than its fPCA counterpart.

These simulated data are intended as an illustrative example rather than as a quantitative investigation of the statistical properties of our method and we do not pursue a simulation here. MAFR by definition reduces the roughness of the leading principal components and can also be expected to reduce their variance. In this particular framework, it is also easy to show that rotating 25 principal components exactly recovers the original fourier basis up to changes of sign and the order of sin and cosin pairs.

4.2 Electricity Demand Data

The Electricity Demand data are obtained from the `R` package `fds` (Shang and Hyndman, 2013). The data comprise the half-hourly demand for electricity in Adelaide, Australia over the period 6/7/1997 to 31/3/2007. Electricity demand in Adelaide is highest in summer and winter. Interestingly the variability in electricity demand is largest when temperature is high. For further information on the Electricity Demand data the reader is referred to Magnano and Boland (2007) and Magnano et al. (2008). In the our analysis we restrict attention to the Monday demand for electricity and consider electricity demand as a function of time of the day. Figure 2 contains plots of the smoothed electricity demand versus time of the day for each of the 508 Mondays over the period of observation.

For this analysis we employed a second derivative rotation to the first five principal component directions – these accounted for 99% of variation in the data.

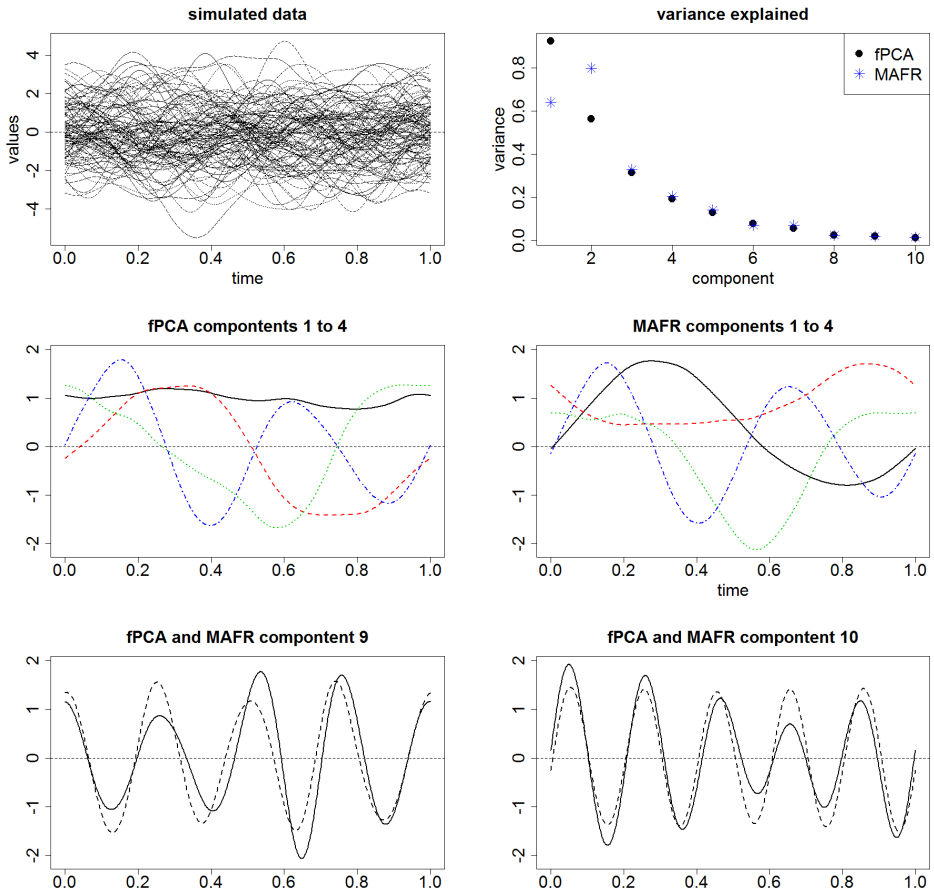

Figure 1: Results of factor rotation based on simulated data with sinusoidal functional principal components. Top row: data (left) and variance components for fPCA (solid circles) and MAFR (stars) components. Second row: leading four fPCA (left) and MAFR (right) components. Bottom row: ninth (left) and tenth (right) fPCA (solid) and MAFR (dashed) components.

The remaining plots in Figure 2 show the fPCA and MAFR components where a smoothing effect is evident, particularly in the second component while the later components largely retain their shape.

5 Discussion

This paper examines the development of factor rotations aimed specifically at providing more interpretable bases for the use of functional principal components. Our approach here has been to find rotations that increase the smoothness of the leading principal components and we find that we are able to provide smoother leading principal components without have a large effect on the more rough components.

The proposed methods are distinct from methods which incorporate smoothing directly into a functional principal components analysis; see, for example Silverman (1996). Here, we fix a subspace on which we will project our data and seek a more interpretable representation of this subspace, rather than rotating the subspace so that the original representation is smoother.

Our methods are also distinct from more classical factor rotation methods in that we target the smoothness of the factors in sequence rather than jointly. A joint rotation criterion could be obtained by considering a weighted sum of smoothing factors:

$$(\mathbf{b}_1, \ldots, \mathbf{b}_k) = \operatorname{argmin} \sum_{j=1}^{k} w_j \int \left[L\mathbf{b}_j^T \boldsymbol{\phi}(t) \right]^2 dt, \text{ such that } \mathbf{b}_i^T \mathbf{b}_j = I_{i=j}$$

which is solved by the eigenvectors of $W^{1/2} P W^{1/2}$ if W is a diagonal matrix with the w_j on the diagonal and for which the MAFR rotation is a limiting case. However, this then poses the problem of how to select the w_j and we argue that even after deciding on a dimension, greatest attention is still paid to the leading components and these should be our focus in factor rotation. Our methods also a variation on those of Liu et al. (2012) in defining correlation with an orthogonal basis with respect to a smoothing norm.

While we have advanced some methods designed at improving the smoothness of principal component functions, we believe that there remains considerable potential for the further development of factor rotations aimed at yielding interpretable bases specifically for functional data.

References

Chiou, J.-M., Müller, H.-G., and Wang, J.-L. (2004), "Functional Response Models," *Statistica Sinica*, 14, 659–677.

Cunningham, J. P. and Ghahramani, Z. (2014), "Unifying linear dimensionality reduction," *arXiv preprint arXiv:1406.0873*.

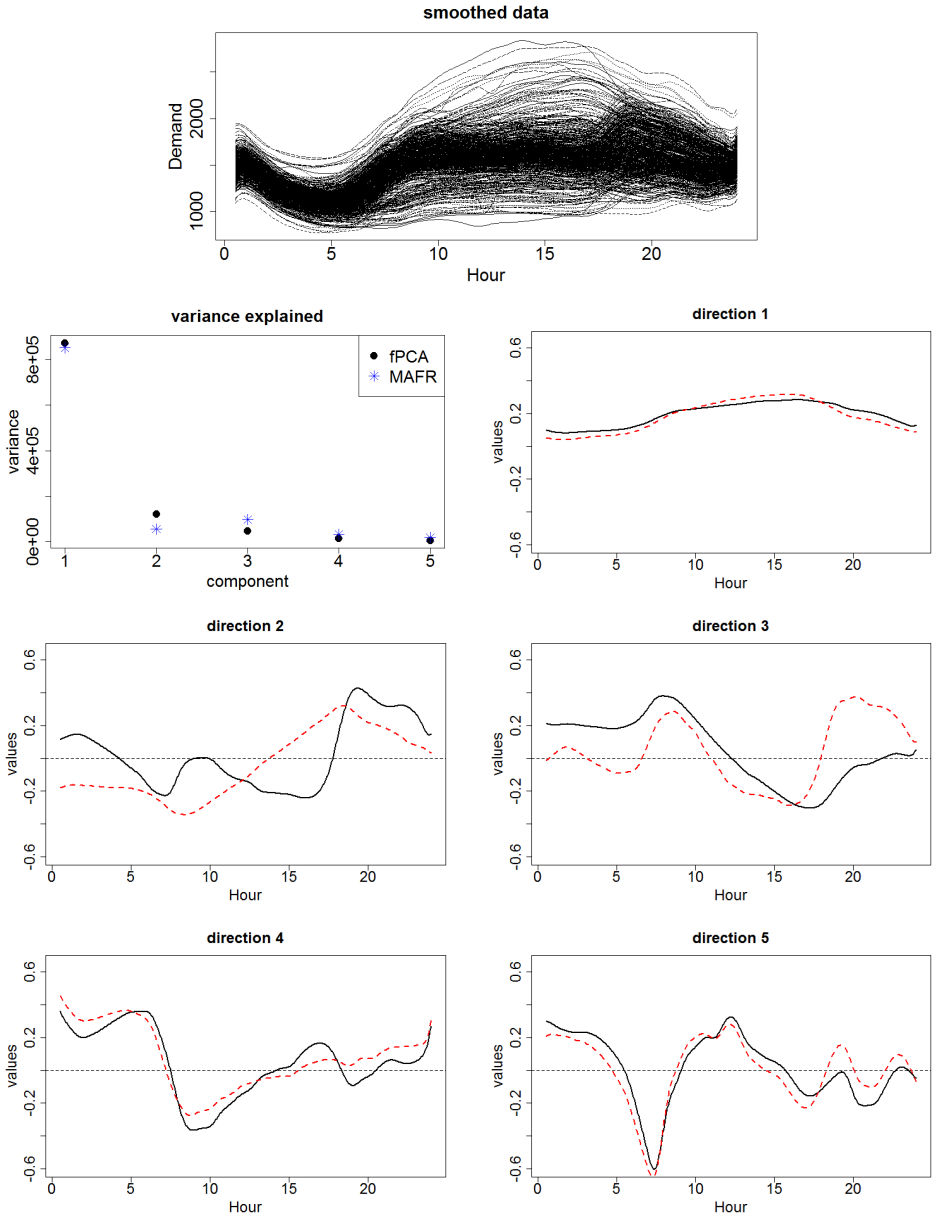

Figure 2: Smoothed Electricity Demand Data (top) along with the fPCA (solid circles) and MAFR (stars) variance components and fPCA (black) and MAFR (red) functions. In this analysis the number of fPCA required to explain 99% of the variation were retained.

Gallagher, N. B., Shaver, J. M., Bishop, R., Roginski, R. T., and Wise, B. M. (2014), "Decompositions using maximum signal factors," *Journal of Chemometrics*, n/a–n/a.

Goldsmith, J., Bobb, J., Crainiceanu, C. M., Caffo, B., and Reich, D. (2011), "Penalized Functional Regression," *Journal of Computational and Graphical Statistics*, 20, 830–851.

Hall, P., Müller, H., and Wang, J. (2006), "Properties of principal component methods for functional and longitudinal data analysis," *The Annals of Statistics*, 34, 1493–1517.

Henderson, A., Fletcher, J. S., and Vickerman, J. C. (2009), "A comparison of PCA and MAF for ToF-SIMS image interpretation," *Surface and Interface Analysis*, 41, 666–674.

Kaiser, H. F. (1958), "The varimax criterion for analytic rotation in factor analysis," *Psychometrika*, 23, 187–200.

Liu, C., Ray, S., Hooker, G., and Friedl, M. (2012), "Functional Factor Analysis For Periodic Remote Sensing Data," *Annals of Applied Statistics*, 6, 601–624.

Magnano, L. and Boland, J. W. (2007), "Generation of synthetic sequences of electricity demand: Application in South Australia," *Energy*, 32, 2230–2243.

Magnano, L., Boland, J. W., and Hyndman, R. (2008), "Generation of synthetic sequences of half-hourly temperature," *Environmetrics*, 19, 818–835.

Müller, H.-G. and Yao, F. (2008), "Functional Additive Models," *Journal of the American Statistical Association*, 103, 1534–1544.

Paul, D. and Peng, J. (2011), "Principal components analysis for sparsely observed correlated functional data using a kernel smoothing approach." *Electron. J. Stat.*, 5, 1960–2003.

Peng, J. and Paul, D. (2009), "A geometric approach to maximum likelihood estimation of the functional principal components from sparse longitudinal data," *Journal of Computational and Graphical Statistics*, 18, 995–1015.

Ramsay, J. O., Hooker, G., and Graves, S. (2009), *Functional Data Analysis in R and Matlab*, New York: Springer.

Ramsay, J. O. and Silverman, B. W. (2005), *Functional Data Analysis, Second Edition*, New York: Springer.

Ramsay, J. O., Wickham, H., Graves, S., and Hooker, G. (2013), *fda: Functional Data Analysis*, r package version 2.4.0.

Rice, J. and Silverman, B. (1991), "Estimating the mean and covariance structure nonparametrically when the data are curves," *Journal of the Royal Statistical Society, Series B*, 53, 233–243.

Sentürk, D. and Müller, H.-G. (2006), "Inference for covariate adjusted regression via varying coefficient models," *The Annals of Statistics*, 34, 654–679.

Shang, H. L. and Hyndman, R. J. (2013), *fds: Functional data sets*, r package version 1.7.

Shapiro, D. and Switzer, P. (1989), "Extracting time trends from multiple monitoring sites: Dept. of Statistics," Tech. rep., Stanford University, Tech. Rep. 132.

Silverman, B. W. (1996), "Smoothed Functional Principal Components by Choice of Norm," *Annals of Statistics*, 24, 1–24.

Switzer, P. and Green, A. (1984), "Min/max autocorrelation factors for multivariate spatial imagery: Dept. of Statistics," Tech. rep., Stanford University, Tech. Rep. 6.

Woillez, M., Rivoirard, J., and Petitgas, P. (2009), "Using min/max autocorrelation factors of survey-based indicators to follow the evolution of fish stocks in time," *Aquatic Living Resources*, 22, 193–200.

Yao, F., Müller, H.-G., and Wang, J.-L. (2005a), "Functional Data Analysis for Sparse Longitudinal Data," *Journal of the American Statistical Association*, 100, 577–590.

— (2005b), "Functional linear regression analysis for longitudinal data," *Annals of Statistics*, 33, 2873–2903.